

Representing Java Classes in a Typed Intermediate Language*

Christopher League Zhong Shao Valery Trifonov
Dept. of Computer Science
Yale University
New Haven, CT 06520
`{league, shao, trifonov}@cs.yale.edu`

Abstract

We propose a conservative extension of the polymorphic lambda calculus (F^ω) as an intermediate language for compiling languages with name-based class and interface hierarchies. Our extension enriches standard F^ω with recursive types, existential types, and row polymorphism, but only ordered records with no subtyping. Basing our language on F^ω makes it also a suitable target for translation from other higher-order languages; this enables the safe interoperation between class-based and higher-order languages and the reuse of common type-directed optimization techniques, compiler back ends, and runtime support.

We present the formal semantics of our intermediate language and illustrate its features by providing a formal translation from a subset of Java, including classes, interfaces, and private instance variables. The translation preserves the name-based hierarchical relation between Java classes and interfaces, and allows access to private instance variables of parameters of the same class as the one defining the method. It also exposes the details of method invocation and instance variable access and allows many standard optimizations to be performed on the object-oriented code.

1 Introduction

The explosive growth of the World Wide Web has, as the Java phenomenon demonstrates, induced newfound interest in mobile computation for “programming the Web.” In this domain, the safety and security properties of programs are more crucial than ever before. Recent work demonstrates a clear connection between security properties and formal type systems [28, 20, 24].

Type checking has long been used to ensure certain properties about the runtime behavior of programs written in strongly typed languages. In the conventional model, however, type information is discarded immediately after type checking. We must therefore *trust* that the compiler—through its many transformations and optimizations—faithfully preserves the source language semantics. Furthermore, given only the object code (without type informa-

* This research was sponsored in part by the Defense Advanced Research Projects Agency ITO under the title “Software Evolution using HOT Language Technology,” DARPA Order No. D888, issued under Contract No. F30602-96-2-0232, and in part by NSF Grant CCR-9633390. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

To appear in ICFP’99: Proceedings of the 1999 ACM SIGPLAN International Conference on Functional Programming, September 1999, Paris, France.

tion), it may be difficult or impossible to verify such properties. Thus, it is becoming increasingly important to preserve full type information throughout the compilation process.

The Java Virtual Machine Language (JVML) [25] was designed to address these issues. JVML byte code contains sufficient type information for an automatic verifier to prove memory safety and other properties. The translation from Java to JVML is mostly type-preserving.

While JVML does contain type information and submits to verification, it has several drawbacks. First, JVML may not be a good fit for source languages other than Java [27, 3]. For example, JVML does not provide direct support for tail-recursion, higher-order functions, and polymorphic functions, making it impractical for implementing functional languages. Second, JVML is designed for bytecode interpretation, not as a compiler intermediate language. The JVML instruction set is based on a stack machine so it is difficult to perform standard optimizations on JVML. Third, JVML has a complex semantics so it does not provide good support for formal reasoning. Much cutting edge research on security and information flow [28, 20, 24] is hard to incorporate into the JVML-based framework.

The FLINT project at Yale [34, 35] takes a different approach. We aim to build a compiler infrastructure for HOT (higher-order and typed) languages. Our goal is to use a richly typed intermediate language (based on the polymorphic λ -calculus F^ω [18, 33]) as a common target for compiling various source languages (e.g., Java, ML). Our system would allow the reuse of common type-directed optimization techniques, compiler back ends, and runtime support.

Building a production-quality type-preserving compiler is by no means trivial, especially in the presence of advanced language features such as parametric polymorphism and higher-order modules. Recent advances in compiler technology, however, are making type-preserving compilation a reality [37, 36]; the current version of FLINT has been in widespread use as the intermediate language of the SML/NJ compiler [2] since January 1997.

Extending our typed intermediate language to handle Java classes poses many new challenges. First, the Java type system is dramatically different from the F^ω type system. Naively combining Java and FLINT could easily lead to an incomprehensible language. The challenge is to abstract the commonality and to find a synergy between the two languages.

Another challenge in modeling Java classes is to find an object encoding that is faithful to the Java semantics yet still supports efficient implementation. Existing encodings of object-oriented features in typed λ -calculi [21, 5] use records of functions as dictionaries and existential types for dynamic binding. These encodings, however, do not support Java-like name-based class hierarchies; two different class types with exactly same set of methods and fields are considered as equivalent by these encodings. This

Report Documentation Page			Form Approved OMB No. 0704-0188	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE SEP 1999	2. REPORT TYPE		3. DATES COVERED -	
4. TITLE AND SUBTITLE Representing Java Classes in a Typed Intermediate Language			5a. CONTRACT NUMBER	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency, 3701 North Fairfax Dr, Arlington, VA, 22203-1714			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT see report				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 14
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified		

is not acceptable because class and interface names in Java play a critical role during type-checking, linking, loading, runtime execution, and casting. The explicit class hierarchy is also crucial for many important optimizations on object-oriented languages [10].

Furthermore, because existing object encodings are intended as theoretical models rather than as intermediate representations, they often use non-trivial coercions in order to simplify the type system. For example, Pierce and Turner [29] model inheritance using a pair of functions for coercing between the superclass and subclass views of an object. For the purposes of compilation, the cost of calling an unknown function to obtain a different view of an object is unacceptable. For Java in particular, these coercions would typically be identity functions. In the presence of separate compilation, we could not expect optimizations to be able to replace these general coercions with more efficient operations.

Finally, the implicit subtyping typically used in these encodings can make them rather complex, in some cases rendering type checking undecidable. We want a more conservative extension, less powerful than, say, F^{ω}_{\leq} [9], but able to support fast type-checking, an explicit class hierarchy, and efficient implementation.

Our main contribution is a formal translation of Java classes, interfaces, and privacy into a variant of F^{ω} using simple and well-known extensions. Although we use several familiar techniques, their combination to support the efficient compilation of an interesting subset of Java is novel.

We use row polymorphism [31] to implement Java inheritance and to allow objects of subclasses to masquerade as superclasses. We use existential types with dot notation [8] to model Java's named types and to preserve the class and interface hierarchy. In addition, we informally describe how to extend our framework to support other interesting features of Java such as dynamic casts, checked exceptions, and mutually recursive declarations.

Although our intermediate language contains full type information, the implementation of object creation, method invocation, and field selection are all quite efficient. This is because the coercions used to implement inheritance and superclass subsumption all reduce to no-ops under a type-erasure dynamic semantics. Although our approach to implementing interfaces appears to be unconventional (casting an object to interface type requires a simple coercion), it allows interface method invocations to have the same cost as ordinary method invocations. From our experience with implementing functional languages, we believe the cost of the coercion can be paid for by the fast method accesses it enables. Since all these operations are implemented in terms of standard type application, record selection, and function application, they are subject to standard optimizations. Furthermore, the interaction of these features with other (non-Java) F^{ω} code can be well understood.

The remainder of this paper formulates the source and target languages (Sections 2 and 3), illustrates the object layout (Section 4), explains the translation algorithm (Section 5), discusses the possible extensions (Section 6), and then closes with related work (Section 7) and conclusions (Section 8).

2 The Source Language

The source language *Javacito* is a small calculus representing a restricted subset of Java. It contains the essence of Java classes, interfaces, and access control, with several simplifications that permit a concise formal semantics and a comprehensible translation.

The program in Figure 1 contains three classes and one interface that demonstrate some interesting features of *Javacito*. Class SPt overrides method move, and inherits methods bump and max. The keyword **super** is used to invoke a statically bound method from the superclass (cf. Figure 1, line 10). Methods are selected

```

1 class Pt {
2   private int x           = 0;
3   public Pt max (Pt other) { (this.x > other.x)? this : other }
4   public void move (int dx) { this.x = this.x + dx; }
5   public void bump ()     { this.move (1); }
6 }
7 interface Zm {public void zoom (int s);}
8 class SPt ext Pt imp Zm {
9   private int scale      = 1;
10  public void move (int dx) { super.move (this.scale * dx); }
11  public void zoom (int s) { this.scale = this.scale * s; }
12 }
13 class Main {
14   private Pt p          = new Pt;
15   private SPt sp        = new SPt;
16   public void zoom2 (Zm z) { z.zoom (2); }
17   public void main () {
18     this.p.bump ();
19     this.zoom2 ((Zm) this.sp);
20     this.sp.bump ();
21     this.zoom2 ((Zm) (SPt) this.p.max ((Pt) this.sp)) // cf. Section 6
22     this.p.max ((Pt) this.sp));
23 }
24 (new Main).main ();

```

Figure 1: Sample *Javacito* program.

$p ::= decl^* exp$	$exp ::= val$
$decl ::= \mathbf{interface} \ i \ \mathbf{ext} \ i^*$	$\{ exp; exp \}^*$
$\{ (msig;)^*\}$	$\mathbf{new} \ c$
$ \ \mathbf{class} \ c \ \mathbf{ext} \ c' \ \mathbf{imp} \ i^*$	$(ty) \ exp$
$\{ field^* meth^* \}$	$exp.m \ (exp^*)$
$field ::= \mathbf{private} \ ty \ f = exp;$	$c \cdot val \triangleright \mathbf{super}.m \ (exp^*)$
$msig ::= \mathbf{public} \ ty \ m \ ((ty \ x)^*)$	$\underline{c \triangleright exp.f}$
$meth ::= msig \ exp$	$\underline{c \triangleright exp.f} = exp$
$ty ::= i \ \ c$	$val ::= x$
$i \in \mathit{InterfaceNames}$	$f \in \mathit{FieldNames}$
$c \in \mathit{ClassNames} \cup \{\mathit{Object}\}$	$m \in \mathit{MethNames}$
$x \in \mathit{VarNames} \cup \{\mathbf{this}\}$	

Figure 2: Syntax of *Javacito*. Underlines indicate annotations required by the contextual operational semantics.

based on the dynamic class of the receiver. When the bump method is invoked on an object of class SPt, it is delegated to the bump method in class Pt (line 5), which then invokes the move method in class SPt (line 10).

Classes have private mutable fields and public methods. A class name may appear recursively in the types of its fields and methods (line 3). A method in class Pt can access the private fields of arguments which statically have the same class type. Objects can be passed to methods expecting arguments of a superclass (line 22) or interface type (line 19).

We also support several features not demonstrated by this sample program. An interface may extend several other interfaces and a class may implement several interfaces. A class may contain a field which is an instance of the class itself (discussed further in Section 5.1.1), and new instances of a class may be created inside its own methods and field initializers.

Java features we do not support in *Javacito* include null refer-

Environment syntax:

Type signature

$$T ::= ty \mid ty_1 \dots ty_n \rightarrow ty$$

Hierarchy

$$\mathcal{H} ::= \emptyset \mid \mathcal{H} \uplus \{i \mapsto [i^*]\} \mid \mathcal{H} \uplus \{c \mapsto \langle c', [i^*] \rangle\}$$

Member type environment

$$\mathcal{E} ::= \emptyset \mid \mathcal{E} \uplus \{ty \mapsto \langle \langle f, ty \rangle^*, \langle m, T \rangle^* \rangle\}$$

Member code environment

$$\mathcal{R} ::= \emptyset \mid \mathcal{R} \uplus \{c \mapsto \langle \langle f, exp \rangle^*, \langle m, (x_1 \dots x_n) exp \rangle^* \rangle\}$$

$\vdash_{\mathcal{H}}$

$$\frac{\vdash_{\mathcal{H}} \mathcal{H} \quad i \notin \text{dom}(\mathcal{H})}{\vdash_{\mathcal{H}} \mathcal{H} \uplus \{i \mapsto [i_1, \dots, i_n]\}}$$

$$\frac{\vdash_{\mathcal{H}} \mathcal{H} \quad c \notin \text{dom}(\mathcal{H}) \quad \mathcal{H} \vdash_t c' \quad \forall j \in \{1..n\}. (i_j = i_{j'} \implies j = j' \text{ and } \mathcal{H} \vdash_t i_j)}{\vdash_{\mathcal{H}} \mathcal{H} \uplus \{c \mapsto \langle c', [i_1, \dots, i_n] \rangle\}}$$

$\vdash_{\mathcal{E}}$

$$\frac{\vdash_{\mathcal{H}} \mathcal{E} \quad \forall j \in \{1..n\}. \mathcal{H} \vdash_t ty_j \quad \forall j \in \{1..n\}. f_j = f_{j'} \implies j = j' \quad \forall k \in \{1..p\}. \mathcal{H} \vdash_t T_k \quad \forall k \in \{1..p\}. m_k = m_{k'} \implies k = k' \quad \forall k \in \{1..p\}. \forall ty'. (ty \prec_{\mathcal{H}} ty' \text{ and } \langle m_k, T' \rangle \in_{\mathcal{H}\mathcal{E}} ty') \implies T_k = T' \quad \forall j \in \{1..n\}. ty \mapsto \langle \langle f_j, ty_j \rangle^{j \leftarrow [1..n]}, \langle m_k, T_k \rangle^{k \leftarrow [1..p]} \rangle \rangle}{\vdash_{\mathcal{H}} \mathcal{E} \uplus \{ty \mapsto \langle \langle f_j, ty_j \rangle^{j \leftarrow [1..n]}, \langle m_k, T_k \rangle^{k \leftarrow [1..p]} \rangle \rangle\}}$$

Figure 3: Environment formation.

ences, public fields, private methods, static members, **protected**, package scope, **final**, constructors, finalizers, mutually recursive classes, exceptions, reflection, and concurrency. As discussed informally in Section 6, many of these features are simple extensions of the framework. Although a dynamic cast appears in the example (line 21), we also treat this feature as an extension.

We use integers, arithmetic operations, conditional expressions, and **void** in the example without defining them formally. We do not distinguish between statements and expressions; rather, we use an ML-like sequence expression with syntax $\{exp; exp\}^*$. The last expression in this sequence is implicitly returned.

Finally, we assume several transformations have been made to the Java code in a pre-processing phase. For instance, all implicit references to **this** are made explicit. Overloaded method references are resolved statically by including argument types as part of the method name. Instance variable shadowing is not an issue in this restricted language because all fields are private, so **super** is only needed for invoking methods in the superclass. Implicit subsumption is made explicit by inserting sequences of upward casts. Similar transformations are defined formally in several papers by Drossopoulou and Eisenbach [11, 12].

The syntax of *Javacito* is given in Figure 2. The underlined terms in Figure 2 are annotations required by our operational semantics. In all terms, the prefix $c \triangleright$ indicates that the term is statically enclosed in the declaration of class c . In the **super** term, the additional val in the annotation holds the value of **this**. These annotations are required by the contextual operational semantics, as explained below.

2.1 Static semantics

Figures 3 and 4 contain the syntax of static environments, rules for environment formation, and definitions of environment-summarizing relations. A hierarchy \mathcal{H} is an environment which maps class or interface names to their immediate superclasses and

Hierarchy relations:

$$c \prec_{\mathcal{H}}^c c' \Leftrightarrow c' = \pi_1(\mathcal{H}(c))$$

Class is declared as an immediate subclass

$$i \prec_{\mathcal{H}}^i i' \Leftrightarrow i' \in \mathcal{H}(i)$$

Interface is declared as an immediate subinterface

$$c \ll_{\mathcal{H}}^c i \Leftrightarrow i \in \pi_2(\mathcal{H}(c))$$

Class declares implementation of an interface

Derived hierarchy relations:

$$\prec_{\mathcal{H}} \equiv \prec_{\mathcal{H}}^c \cup \prec_{\mathcal{H}}^i \cup \ll_{\mathcal{H}}^c: \text{Type is an immediate subtype}$$

$$\leq_{\mathcal{H}}^c \equiv \text{reflexive transitive closure of } \prec_{\mathcal{H}}^c: \text{Class is a subclass}$$

$$\leq_{\mathcal{H}}^i \equiv \text{reflexive transitive closure of } \prec_{\mathcal{H}}^i: \text{Interface is a subinterface}$$

$$c \ll_{\mathcal{H}}^c i \Leftrightarrow \exists c', i' \text{ s.t. } c \leq_{\mathcal{H}}^c c' \text{ and } c' \ll_{\mathcal{H}}^c i' \text{ and } i' \leq_{\mathcal{H}}^i i$$

Class implements an interface

$$\leq_{\mathcal{H}} \equiv \leq_{\mathcal{H}}^c \cup \leq_{\mathcal{H}}^i \cup \ll_{\mathcal{H}}^c: \text{Type is a subtype}$$

Membership relations:

$$\langle f, X \rangle \in_{\mathcal{E}}^c c \Leftrightarrow \langle f, X \rangle \in \pi_1(\mathcal{E}(c)): \text{Field is declared in class}$$

$$\langle m, X \rangle \in_{\mathcal{E}}^t ty \Leftrightarrow \langle m, X \rangle \in \pi_2(\mathcal{E}(ty)): \text{Method is declared in type}$$

Derived membership relations:

$$\langle f, X \rangle \in_{\mathcal{H}\mathcal{E}}^c c \Leftrightarrow \langle f, X \rangle \in_{\mathcal{E}}^c c' \text{ and } c \leq_{\mathcal{H}}^c c'$$

Field is contained in class

$$\langle m, X \rangle \in_{\mathcal{H}\mathcal{E}}^c c \Leftrightarrow \langle m, X \rangle \in_{\mathcal{E}}^c c' \text{ and }$$

$$c' = \min\{c'' \mid c \leq_{\mathcal{H}}^c c'' \text{ and } \langle m, _ \rangle \in_{\mathcal{E}}^c c''\}$$

Method is contained in class

$$\langle m, X \rangle \in_{\mathcal{H}\mathcal{E}}^i i \Leftrightarrow \langle m, X \rangle \in_{\mathcal{E}}^i i' \text{ and } i \leq_{\mathcal{H}}^i i'$$

Method is contained in interface

$$\in_{\mathcal{H}\mathcal{E}} \equiv \in_{\mathcal{H}\mathcal{E}}^c \cup \in_{\mathcal{H}\mathcal{E}}^i: \text{Member is contained in type}$$

Figure 4: Environment relations.

superinterfaces. The hierarchy formation judgment $\vdash_{\mathcal{H}}$ guarantees that a new type is added to the hierarchy at most once, and that all supertypes are already defined and mentioned only once. The member type environment \mathcal{E} maps types to lists of field and methods names, paired with the respective types. The formation rule ensures that the names are unique, the types are valid, and that overriding a method preserves its type.¹ The member code environment \mathcal{R} maps class names to a set of field initializers and method bodies. This environment is created during elaboration, for use at runtime.

We use a set of hierarchy and membership relations similar to those used to describe CLASSICJAVA [17], except that ours are defined in terms of environments \mathcal{H} and \mathcal{E} rather than on complete programs. The membership relations apply to both type and code environments. Note that for methods contained in class c ($\langle m, X \rangle \in_{\mathcal{H}\mathcal{E}}^c c$), the auxiliary datum X is drawn from the *nearest* superclass of c that declares the method m . When a code environment \mathcal{R} is used in place of \mathcal{E} , this has the effect of selecting the appropriate method body by using the inheritance hierarchy.

Figure 5 contains the typing judgments comprising the static semantics of *Javacito*. Declarations are processed one at a time; they may be recursive, but not mutually recursive. All declarations extend the hierarchy \mathcal{H} and the member type environment \mathcal{E} . Class declarations also extend the member code environment \mathcal{R} , which is used to retrieve code at runtime. Unlike in CLASSICJAVA, there are no rules for subsumption. Instead, we require explicit upward casts. Casts are only legal from some type to its immediate supertype (exp-cast). Privacy is enforced in (exp-get) and (exp-set), the rules for field access. A field f can only be accessed from within the class c that declares it. The object from which f is selected may be

¹Rules for correctness of empty environments are omitted. Our notation for lists and list comprehensions is the usual, but we abuse set membership notation on lists.

\vdash_p <p>(prog-decl)</p> $\frac{\mathcal{H}; \mathcal{E}; \mathcal{R} \vdash_d \text{decl} \Rightarrow \mathcal{H}'; \mathcal{E}'; \mathcal{R}'}{\mathcal{H}'; \mathcal{E}'; \mathcal{R}' \vdash_p \text{decl}_1 \dots \text{decl}_n \text{ exp} : \text{ty} \Rightarrow \mathcal{H}''; \mathcal{R}''}$ $\mathcal{H}; \mathcal{E}; \mathcal{R} \vdash_p \text{decl decl}_1 \dots \text{decl}_n \text{ exp} : \text{ty} \Rightarrow \mathcal{H}'''; \mathcal{R}'''$	\vdash_d <p>(decl-interface)</p> $\frac{\mathcal{H}' = \mathcal{H} \uplus \{i \mapsto [i_1, \dots, i_n]\} \quad \mathcal{E}' = \mathcal{E} \uplus \{i \mapsto \langle[], [\text{MT}(msig_k)^{k \leftarrow [1..m]}]\rangle\} \quad \vdash_H \mathcal{H}' \quad \mathcal{H}' \vdash_E \mathcal{E}'}{\mathcal{H}; \mathcal{E}; \mathcal{R} \vdash_d \text{interface } i \text{ ext } i_1 \dots i_n \{msig_1; \dots msig_m;\} \Rightarrow \mathcal{H}'; \mathcal{E}'; \mathcal{R}}$
\vdash_t <p>(typ-ok)</p> $\frac{ty \in \text{dom}(\mathcal{H}) \cup \{\text{Object}\}}{\mathcal{H} \vdash_t ty}$	<p>(typ-msig)</p> $\frac{\forall j \in \{1..n\}. \mathcal{H} \vdash_t ty_j}{\mathcal{H} \vdash_t ty}$ $\frac{\mathcal{H} \vdash_t ty}{\mathcal{H} \vdash_t ty_1 \dots ty_n \rightarrow ty}$
\vdash_m <p>(meth-body)</p> $\frac{\mathcal{H}; \mathcal{E}; \{\text{this} : c, x_1 : ty_1, \dots, x_n : ty_n\}; c \vdash_e \text{exp} : ty}{\mathcal{H}; \mathcal{E}; c \vdash_m ty m(ty_1 x_1, \dots, ty_n x_n) \text{ exp}}$	$\frac{\mathcal{H}' = \mathcal{H} \uplus \{c \mapsto \langle c', [i_1 \dots i_n]\rangle\} \quad \vdash_H \mathcal{H}' \quad \mathcal{E}' = \mathcal{E} \uplus \{c \mapsto \langle \langle f_k, ty_k \rangle^{k \leftarrow [1..m]}, [\text{MT}(msig_l)^{l \leftarrow [1..p]}\rangle\}\} \quad \vdash_E \mathcal{E}' \quad \mathcal{R}' = \mathcal{R} \uplus \{c \mapsto \langle \langle f_k, exp_k \rangle k \in \{1..m\}, \{\text{MR}(msig_l exp'_l) l \in \{1..p\}\}\rangle\} \quad \forall j \in \{1..n\}. \text{ if } \langle m', T' \rangle \in_{\mathcal{H} \mathcal{E} j}^i \text{ then } \langle m', T' \rangle \in_{\mathcal{H} \mathcal{E} c}^e \quad \forall k \in \{1..m\}. \mathcal{H}'; \mathcal{E}'; \emptyset; c \vdash_e exp_k : ty_k \quad \forall l \in \{1..p\}. \mathcal{H}'; \mathcal{E}'; c \vdash_m msig_l exp'_l}{\mathcal{H}; \mathcal{E}; \mathcal{R} \vdash_d \{ty_1 f_1 = exp_1; \dots ty_m f_m = exp_m; \dots msig_1 exp'_1 \dots msig_p exp'_p\} \Rightarrow \mathcal{H}'; \mathcal{E}'; \mathcal{R}'}$
\vdash_e <p>(exp-var)</p> $\frac{x \in \text{dom}(\Gamma)}{\vdash_e x : \Gamma(x)}$	<p>(exp-cast)</p> $\frac{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e \text{exp} : ty' \quad ty' \prec_H ty}{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e (ty) \text{ exp} : ty}$
<p>(exp-call)</p> $\frac{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e \text{exp} : ty' \quad \langle m, (ty_1 \dots ty_n \rightarrow ty) \rangle \in_{\mathcal{H} \mathcal{E} t} ty' \quad \forall j \in \{1..n\}. \mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e exp_j : ty_j}{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e \text{exp.m(exp}_1 \dots \text{exp}_n\text{)} : ty}$	<p>(exp-new)</p> $\frac{\mathcal{H} \vdash_t c}{\mathcal{H}; \vdash_e \text{new } c : c}$
<p>(exp-super)</p> $\frac{c \prec_H^C c' \quad \langle m, (ty_1 \dots ty_n \rightarrow ty) \rangle \in_{\mathcal{H} \mathcal{E} t} c' \quad \forall j \in \{1..n\}. \mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e exp_j : ty_j}{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e c \cdot \text{this} \triangleright \text{super}.m(\text{exp}_1 \dots \text{exp}_n) : ty}$	<p>(exp-body)</p> $\frac{\forall j \in \{1..n\}. \mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e exp_j : ty_j}{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e \{exp_1; \dots; exp_n\} : ty_n}$
<p>(exp-get)</p> $\frac{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e \text{exp} : ty' \quad ty' \leq_H^C c \quad \langle f, ty \rangle \in_{\mathcal{E}}^C c}{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e \underline{c \triangleright exp.f} : ty}$	<p>(exp-set)</p> $\frac{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e exp_1 : ty' \quad ty' \leq_H^C c \quad \langle f, ty \rangle \in_{\mathcal{E}}^C c}{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash_e \underline{c \triangleright exp_1.f} = exp_2 : ty}$

Figure 5: Static semantics of Javacito.

a subclass of c , however.

Proposition 1 (Decidability) \vdash_p is decidable.

2.2 Operational semantics

Figure 6 contains the contextual small-step operational semantics of Javacito. We introduce two new terms: object expressions $\langle c, [(f \mapsto \text{exp})^*]\rangle$ and locations loc . The runtime store \mathcal{S} maps each location to an object value, which consists of a class tag and a collection \mathcal{F} of field-value bindings.

An evaluation context E is an expression with a hole, which encodes the search for the next redex. $\{E\dots\}$ indicates that the first expression of a sequence is always the redex. An argument list of the form $(val_1, \dots, val_n, E, \dots)$ means that all expressions to the left of the redex must be values. The state of the computation is characterized completely by a term paired with a store. The primitive reduction relation \hookrightarrow is defined between these states, given the class hierarchy \mathcal{H} and the code \mathcal{R} . We do not need error states; since Javacito does not have null references or dynamic casts, no runtime errors can occur. Rather, evaluation gets stuck if one of the side conditions of a reduction is not met.

An object expression is used to represent partially initialized objects during the reduction of $\text{new } c$. Once all the field expressions are reduced to values, we allocate a new location in which to store the object literal.

Note how the annotations are used in defining reductions. In (r-super), we retrieve the method m from c' , the immediate superclass of the class c which contains the **super** call. The *loc* annotation in (r-super) is the location of **this**. In the rules (r-get) and (r-set) for field access, the class annotation is used to enforce privacy; f must be declared in the class which contains the field access expression.

Even though an upward cast has no effect at runtime, the side condition of (r-cast) ensures that evaluation will get stuck if a cast is somehow applied to an object of inappropriate type.

Javacito is sound, meaning that a program which passes the static type checking will not get stuck at runtime. Thanks to the strong side conditions on (r-get) and (r-set), this also means that privacy will not be violated at runtime. To formally state and prove soundness we use standard techniques and extend the typing relation to a relation $\mathcal{H}; \mathcal{R} \vdash_c \langle \text{exp}, \mathcal{S} \rangle : ty$ on runtime configurations (see e.g. [13]), which allows us to prove the subject reduction and progress properties. Details can be found in the accompanying technical report[23].

3 The Target Language

The target language MINIFLINT is a call-by-value variant of the omega-order polymorphic lambda calculus (F^ω) [18, 33]. We extend F^ω with standard constructs for recursive types of base kind (with explicit **fold** and **unfold**), existential types with dot nota-

Syntactic extensions:

Object expression $exp ::= \dots | \langle c, [(f \mapsto exp)^*] \rangle$
 Location value $val ::= \dots | loc$

Field map $\mathcal{F} ::= \emptyset | \mathcal{F} \sqcup [f \mapsto val]$
 Runtime store $\mathcal{S} ::= \emptyset | \mathcal{S} \sqcup \{loc \mapsto \langle c, \mathcal{F} \rangle\}$

Evaluation contexts:

$$E ::= \bullet | \{ E \dots \} | (ty) E$$

$$\quad | \langle c, [f \mapsto val_1, \dots, f \mapsto val_n, f \mapsto E, \dots] \rangle$$

$$\quad | E.m(\dots) | val.m(val_1, \dots, val_n, E, \dots)$$

$$\quad | \underline{c \cdot val} \triangleright \underline{\text{super}.m}(val, \dots, val_n, E, \dots)$$

$$\quad | \underline{c \triangleright E.f} | \underline{c \triangleright E.f = exp} | \underline{c \triangleright val.f = E}$$

Primitive reductions:

- (r-new) $\mathcal{H}; \mathcal{R} \vdash \langle E[\mathbf{new} c], \mathcal{S} \rangle \hookrightarrow \langle E[\langle c, \mathcal{F} \rangle], \mathcal{S} \rangle$
 where $\mathcal{F} = [f \mapsto exp \mid \langle f, exp \rangle \in_{\mathcal{HR}}^c c]$
- (r-alloc) $\mathcal{H}; \mathcal{R} \vdash \langle E[\langle c, [f_j \mapsto val_j]^{j \in \{1..n\}} \rangle \mapsto], \mathcal{S} \rangle$
 $\hookrightarrow \langle E[loc], \mathcal{S}[loc \mapsto \langle c, [f_j \mapsto val_j]^{j \in \{1..n\}} \rangle] \rangle$
 where $loc \notin \text{dom}(\mathcal{S})$
- (r-body) $\mathcal{H}; \mathcal{R} \vdash \langle E[\{val; exp_1; \dots; exp_n\}], \mathcal{S} \rangle$
 $\hookrightarrow \langle E[\{exp_1; \dots; exp_n\}], \mathcal{S} \rangle$
- (r-return) $\mathcal{H}; \mathcal{R} \vdash \langle E[\{val\}], \mathcal{S} \rangle \hookrightarrow \langle E[val], \mathcal{S} \rangle$
- (r-call) $\mathcal{H}; \mathcal{R} \vdash \langle E[loc.m(val_1 \dots val_n)], \mathcal{S} \rangle$
 $\hookrightarrow \langle E[exp[loc/\mathbf{this}, val_1/x_1 \dots val_n/x_n]], \mathcal{S} \rangle$
 where $\mathcal{S}(loc) = \langle c', _ \rangle$
 and $\langle m, (x_1 \dots x_n) exp \rangle \in_{\mathcal{HR}}^c c'$
- (r-super) $\mathcal{H}; \mathcal{R} \vdash \langle E[c \cdot loc \triangleright \underline{\text{super}.m}(val_1 \dots val_n)], \mathcal{S} \rangle$
 $\hookrightarrow \langle E[exp[loc/\mathbf{this}, val_1/x_1 \dots val_n/x_n]], \mathcal{S} \rangle$
 where $c \prec_{\mathcal{H}}^c c'$
 and $\langle m, (x_1 \dots x_n) exp \rangle \in_{\mathcal{HR}}^c c'$
- (r-get) $\mathcal{H}; \mathcal{R} \vdash \langle E[\underline{c \triangleright loc.f}], \mathcal{S} \rangle \hookrightarrow \langle E[val], \mathcal{S} \rangle$
 where $\mathcal{S}(loc) = \langle c', \mathcal{F} \rangle$ and $\mathcal{F}(f) = val$
 and $\langle f, _ \rangle \in_{\mathcal{R}}^c c$ and $c' \leq_{\mathcal{H}}^c c$
- (r-set) $\mathcal{H}; \mathcal{R} \vdash \langle E[\underline{c \triangleright loc.f = val}], \mathcal{S} \rangle$
 $\hookrightarrow \langle E[val], \mathcal{S}[loc \mapsto \langle c, \mathcal{F}' \rangle] \rangle$
 where $\mathcal{S}(loc) = \langle c', \mathcal{F} \rangle$ and $c' \leq_{\mathcal{H}}^c c$
 and $\mathcal{F}' = \mathcal{F}[f \mapsto val]$ and $\langle f, _ \rangle \in_{\mathcal{R}}^c c$
- (r-cast) $\mathcal{H}; \mathcal{R} \vdash \langle E[(ty) loc], \mathcal{S} \rangle \hookrightarrow \langle E[loc], \mathcal{S} \rangle$
 where $\mathcal{S}(loc) = \langle c, _ \rangle$ and $c \leq_{\mathcal{H}} ty$

Figure 6: Operational semantics of Javacito.

tion [8], tuples with row polymorphism, and a store. The syntax is given in Figure 7, and the static semantics in Figure 8. Notice that we only list the formation rules for environments, types, and terms; the type reduction and equivalence rules are omitted.

Following Rémy [31] we introduce a kind of rows R^L , where L is the set of labels banned from the row. \mathbf{Abs}^L is the empty row missing L , and $l : \tau; \tau'$ constructs a row from the element l of type τ and the row τ' . The record constructor $\{ \cdot \}$ lifts a complete row type (with no labels missing) to kind Ω .

Consider the following example, a function which fetches field l from a record, regardless of what other fields are present:

```
let fst =  $\Lambda t::R^{\{l\}}. \lambda x : \{l : \mathbf{Int}; t\}. (x.l)$ 
in fst [ $\mathbf{Abs}^{\{l\}}$ ] {l = 5} + fst [m : Int;  $\mathbf{Abs}^{\{l,m\}}$ ] {l = 6, m = 7}
```

We say the function `fst` is polymorphic in the *tail* of its argument.

$$\begin{aligned} \kappa &::= \Omega \mid R^L \mid \kappa \rightarrow \kappa' \\ \tau &::= \mathbf{Ref} \tau \mid \tau \rightarrow \tau' \\ &\quad | t \mid x.\mathbf{Typ} \mid \forall t::\kappa. \tau \mid \exists t::\kappa. \tau \mid \mu t. \tau \mid \lambda t::\kappa. \tau \mid \tau \tau' \\ &\quad | \mathbf{Abs}^L \mid l : \tau; \tau' \mid \{\tau\} \\ e &::= \mathbf{ref} e \mid !e \mid e := e' \\ &\quad | x \mid \lambda x : \tau. e \mid e x \mid \mathbf{let} x = e' \mathbf{in} e \\ &\quad | \mathbf{fold} e \mathbf{as} \tau \mid \mathbf{unfold} e \\ &\quad | \Lambda t::\kappa. e \mid e[\tau] \mid \langle t::\kappa = \tau, e : \tau' \rangle \mid x.\mathbf{val} \\ &\quad | \{l = e\} \mid e.l \end{aligned}$$

$$t \in TypeVars \quad x \in Vars \quad l \in Labels \quad L \in \mathcal{P}(Labels)$$

Derived forms:

$$\{l_1 : \tau_1; \dots; l_n : \tau_n\} \triangleq \{l_1 : \tau_1; \dots; l_n : \tau_n; \mathbf{Abs}^{\{l_1, \dots, l_n\}}\}$$

$$\mathbf{1} \triangleq \{\mathbf{Abs}^\emptyset\}$$

$$e e' \triangleq \mathbf{let} x = e \mathbf{in} \mathbf{let} x' = e' \mathbf{in} x x',$$

where x is not free in e'

Figure 7: Syntax of MINIFLINT.

Types

$$\frac{\vdash A \mathbf{env} \quad A(x) = \exists t::\kappa. \tau}{A \vdash x.\mathbf{Typ} :: \kappa} \quad \frac{\vdash A \mathbf{env}}{A \vdash \mathbf{Abs}^L :: R^L}$$

$$\frac{A \vdash \tau :: \Omega \quad A \vdash \tau' :: R^{L \cup \{l\}}}{A \vdash l : \tau; \tau' :: R^{L - \{l\}}} \quad \frac{A \vdash \tau :: R^\emptyset}{A \vdash \{\tau\} :: \Omega}$$

Terms

$$\frac{A, x : \tau' \vdash e : \tau \quad A \vdash \tau :: \Omega \quad A \vdash e : \{l_1 : \tau_1; \dots; l_n : \tau_n; \tau'\}}{A \vdash \lambda x : \tau'. e : \tau' \rightarrow \tau} \quad \frac{}{A \vdash e.l_n : \tau_n}$$

$$\frac{\forall j \in \{1..n\}. A \vdash e_j : \tau_j}{A \vdash \{l_1 = e_1, \dots, l_n = e_n\} : \{l_1 : \tau_1; \dots; l_n : \tau_n\}}$$

$$\frac{A \vdash e : \tau[\mu t. \tau/t]}{A \vdash \mathbf{fold} e \mathbf{as} \mu t. \tau : \mu t. \tau} \quad \frac{A \vdash \tau' :: \kappa \quad A \vdash e : \tau[\tau'/t]}{A \vdash \langle t::\kappa = \tau', e : \tau \rangle : \exists t::\kappa. \tau}$$

$$\frac{A \vdash e : \mu t. \tau}{A \vdash \mathbf{unfold} e : \tau[\mu t. \tau/t]} \quad \frac{\vdash A \mathbf{env} \quad A(x) = \exists t::\kappa. \tau}{A \vdash x.\mathbf{val} :: \tau[x.\mathbf{Typ}/t]}$$

Figure 8: Static semantics of MINIFLINT—less common rules.

Since these are ordered records, the use of labels here is strictly for clarity of presentation. In an implementation L could be an integer indicating the offset of the row from the beginning of the enclosing record. Note also that record terms are just non-extensible tuples.

An existential type $\exists t::\kappa. \tau$ is the type of a pair containing a hidden type τ' of kind κ and a term of type $\tau[\tau'/t]$. We use dot notation [8] to access the type $(x.\mathbf{Typ})$ and term $(x.\mathbf{val})$ components of a package x . Although this is a form of dependent type, the calculus remains decidable because we restrict the \mathbf{Typ} operation to term variables x . We provide a \mathbf{let} construct to bind term variables and limit their scope. Abstract types $x.\mathbf{Typ}$ and $x'.\mathbf{Typ}$ are equivalent if and only if x and x' are bound by the same \mathbf{let} . The argument of an application must be a variable so that we can track this equivalence through function applications. Variables are prevented from escaping the scope of λ -abstractions $\lambda x : \tau. e$ by banning free occurrences of x in the type of e .

Proposition 2 (Decidability) *The MINIFLINT typing relation $A \vdash e : \tau$ is decidable.*

We omit the dynamic semantics of MINIFLINT, which defines

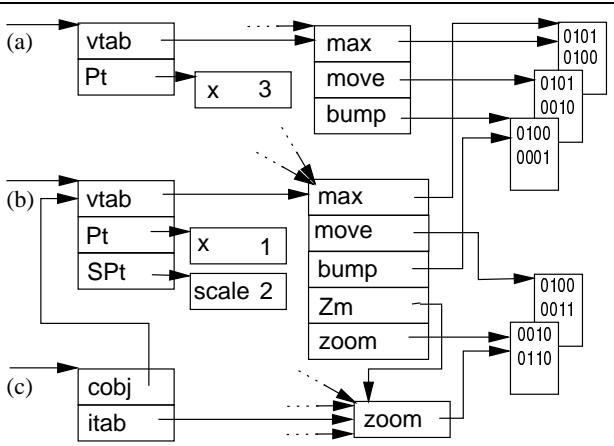


Figure 9: Object layout. (a) and (b) are instances of classes Pt and SPt, respectively. (c) is a view of (b) through the Zm interface.

locations, stores S , values v , and a reduction relation \leftrightarrow between term-store pairs. Since runtime behavior does not depend on types, we can use a type-erasure semantics; this implies that **fold**, **unfold**, type application, package creation, and package projection are all no-ops at runtime. However for proving type soundness and correctness of the translation of *Javacito* to MINIFLINT we define a reduction semantics on typed terms. The reader is referred to the accompanying technical report for details.

4 Object Layout

We start by examining the object layout used in our encoding. Figure 9 depicts the runtime representation of two Java class objects and an interface object. The classes and interfaces are taken from the example in Figure 1.

An object is a record containing a virtual table (*vtable*) and the private fields. The vtable is a per-class data structure containing pointers to the methods defined in the class or inherited from the superclass.

The object marked (a) is an instance of the Pt class, which has one field and three methods. The object marked (b) is an instance of the subclass SPt, which extends Pt with a scale field and a zoom method. SPt overrides the move method; the others are inherited from the superclass by simply pointing to the same code. Note that both the object record and the vtable of (a) are prefixes of those in (b). This property allows superclass methods to access objects created by subclasses without any coercions.

In addition to its four methods, the vtable for class SPt contains a pointer to an interface table (*itable*). An itable is merely a rearrangement of some of the methods in the vtable to correspond to an interface which the class implements. Given an object of interface type, we know nothing about the shape of its vtable. Thus, we use the itable to give a consistent view of the methods in a particular interface, independent of the underlying vtable.

There are various ways of locating methods in interface objects. One technique, used in the Toba ahead-of-time compiler [30], is to construct a per-class dictionary that maps method names to offsets in the vtable. Another idea, implemented in the CACAO 64-bit JIT compiler [22] and described informally elsewhere [26], is to construct an itable for each interface we implement and store them all somewhere in the vtable. Although [22] is not clear on how to use the itable, there appear to be two choices. First, we can search for the appropriate itable in the vtable, which amounts to dictionary

lookup on interface names rather than on method names. Second, when casting an object from class type to interface type, we can select the itable and then pair it with the object itself. This avoids name lookup entirely but requires minor coercions when casting to and between interface types. The object marked (c) in Figure 9 is the result of casting an SPt object (b) to the Zm interface.

In our encoding, we use coercions to achieve fixed-offset method selection regardless whether we are using a class or interface object. From our intuition and experience with implementing functional languages, we believe the cost of the coercion will be paid for by the fast method accesses it enables. Another reason to use the coercion in a typed implementation is that the dictionary lookup (and caching the results for quicker access later) is extremely difficult to prove type-safe.

5 Translation

The main result of this paper is a type-preserving translation algorithm that compiles a well-typed *Javacito* program into a term of the target language MINIFLINT. In this section, we first informally describe our algorithm, using specific examples and a set of type and term macros (see Figures 10 and 11). We then give the formal algorithm (see Figure 12) and show that it is both type-correct and sound.

Each class declaration is separately translated into a let-bound existential package. A skeleton of this package for a class c follows:

```
let  $x_c = \langle t_{own}::\Omega = \boxed{5.1.1},$ 
    {dict =  $\Lambda t_{subFds}.\Lambda t_{subVtab}.\{ \boxed{5.3} \}$  },
    init =  $\lambda : 1. \{ \boxed{5.1.3} \}$  ,
    new =  $\lambda : 1. \boxed{5.3.1}$  ,
    } :  $\boxed{5.1.2, 5.2} \rangle$ 
in ...
```

Missing components in the package have been replaced with boxes indicating the subsections below which contain further detail.

The package is bound to x_c , a variable obtained directly from the class name c . Formally, x_c is a map from *ClassNames* to *Vars*. In examples, we distinguish the two domains using different fonts; the package for class Pt is bound to the variable Pt.

The type component (t_{own}) of the package contains the type of the private fields declared in class c . Outside the package, this type can be selected using dot notation ($x_c.\textbf{Typ}$) but it remains abstract. The value part of the package is a record containing an extensible dictionary (dict), a field initializer (init), and a constructor (new). We say dict is *extensible* because it is parameterized by row types for additional fields and methods. To produce a vtable from an extensible dictionary, we instantiate these row types to empty. Subclasses will instantiate them according to their own new fields and methods.

In Section 5.1, we describe the encoding of private fields and field initialization. Section 5.2 covers the typing issues for vtables and objects. Terms for method invocation, **super**, **new**, and casts are given in Section 5.3. Inheritance and interfaces are handled in sections 5.4 and 5.5, respectively. Finally, Section 5.6 makes some formal claims about the translation.

5.1 Fields and privacy

5.1.1 Own fields

Let the fields declared in class c be represented by a record of ML-style polymorphic **ref** cells called *OwnFlds*[c]. For example,

$OwnFlds[Pt] = \{x : \text{Ref Int}\}$. Using a more efficient flat mutable record presents no problems, but for simplicity we prefer to preserve the orthogonality of features in MINIFLINT.

A tricky issue in defining $OwnFlds[c]$ is that the class c may contain a field of type c , whose privates are also visible to code within class c . Imagine a class List which has fields ‘int’ data’ and ‘List next’. Within class List, it is legal to access not only ‘this.next’ but also ‘this.next.next’ and ‘this.next.next.next’, and so on, ad infinitum.

This implies that, in the general case, $OwnFlds[c]$ must be a recursive type:

$$OwnFlds[c] = \mu t_{\text{own}}. OwnFldsGen[c] (ObjGen[c] t_{\text{own}})$$

where $ObjGen[c]$ (see Section 5.2.2) is the type of a class object parameterized by the type of its private fields. (All type operators used in the translation are listed in Figure 10.) $OwnFldsGen[c]$ takes an argument t_{twins} , creates a record containing **ref** cells for each field declared in c , and makes any fields of type c have type t_{twins} instead. For the class List mentioned above,

$$OwnFldsGen[\text{List}] = \lambda t_{\text{twins}}. \{ \text{data} : \text{Ref Int}; \text{next} : \text{Ref } t_{\text{twins}} \}$$

Substituting this into $OwnFlds[\text{List}]$, we have:

$$OwnFlds[\text{List}] = \mu t_{\text{own}}. \{ \text{data} : \text{Ref Int}; \\ \text{next} : \text{Ref } (ObjGen[\text{List}] t_{\text{own}}) \}$$

As required, the private fields of List include a next field which, in turn, is an object whose private fields include a next field, etc.

5.1.2 All fields

An object contains not only the fields of its own class, but the fields of all its superclasses as well. Field visibility, however, is based on the class in which the fields were declared. Fields from each class can be hidden or revealed independently. This implies that we need to segregate the fields by the class in which they were declared.

Let $AllFldsGen[c]$ produce the type of the segregated record containing all fields in an object of class c . In the example,

$$\begin{aligned} AllFldsGen_f[Pt] &= \lambda t_{\text{own}}. \{ \text{Pt} : t_{\text{own}} \} \\ AllFldsGen_f[SPT] &= \lambda t_{\text{own}}. \{ \text{Pt} : \text{Pt.Typ}; \text{SPT} : t_{\text{own}} \} \end{aligned}$$

where Pt.Typ is the abstract type of private fields in Pt. t_{own} can be instantiated to $OwnFlds[c]$ for use inside the class c , or to $x_c.\text{Typ}$ outside, where x_c is bound to the package implementing class c .

We use row polymorphism to extend superclass field record types in subclasses. We additionally parameterize the above operators by t_{subFlds} , a row of field records defined by potential subclasses. Now we can define the SPT fields using the Pt fields directly:

$$\begin{aligned} AllFldsGen[Pt] &= \lambda t_{\text{own}}. \lambda t_{\text{subFlds}}. \{ \text{Pt} : t_{\text{own}}; t_{\text{subFlds}} \} \\ AllFldsGen[SPT] &= \lambda t_{\text{own}}. \lambda t_{\text{subFlds}}. \\ &\quad AllFldsGen[\text{Pt}] \text{ Pt.Typ } (\text{SPT} : t_{\text{own}}; t_{\text{subFlds}}) \end{aligned}$$

In the general case, the segregated record type is produced using two mutually recursive operators (see Figure 10).

5.1.3 Field initialization

As a part of the code for class c we define a function init which creates a record of **ref** cells, each initialized with the translated field initializer expression. Every class x_c exports init as a function of type $\mathbf{1} \rightarrow x_c.\text{Typ}$. For the Pt class in our example, $Pt.\text{val}.init$ is

```

UnpackObj[c] (town, o : ObjGen[c] town, •) =
let x = unfold o in let x' = x.val in •
PackObj[c]
(town, tsubFlds, tsubVtab, self : SelfGen[c] town tsubFlds tsubVtab) =
fold (tsubFlds = tsubFlds,
      tsubVtab = tsubVtab,
      self : SelfGen[c] town tsubFlds tsubVtab)
      : ∃tsubVtab. SelfGen[c] town tsubFlds tsubVtab)
as ObjGen[c] town
UnpackView[i] (io : ExtView[i], •) = let x = unfold io in •
PackView[i] (tcobj, o : tcobj, it : ExtItab[i] tcobj) =
fold (tcobj::Ω = tcobj,
      {cobj = o, itab = it}
      : ViewTmpl[i] (xi.Typ) ExtView[i] tcobj)
as ExtView[i]
GetFields[c] (o : ObjGen[c] OwnFlds[c]) =
UnpackObj[c]
(OwnFlds[c], o, unfold (unfold (x'.val).fields.lc))
Call[c] (town, o : ObjGen[c] town, m) =
UnpackObj[c] (town, o, unfold (x'.val).vtab.lm (x'.val))
Call[i] (io, io : ExtView[i], m) =
UnpackView[i] (io, x.val.itab.lm (x.val.cobj))
Cast[c, c'] (town, t'own o : ObjGen[c] town) =
UnpackObj[c] (town, o,
               PackObj[c'] (t'own, (lc : town; x.Typ),
                           NewPublic[c] ExtObj[c] x'.Typ, x'.val))
if c' = Super[c]
Cast[c, i] (town, io : ObjGen[c] town) =
UnpackObj[c] (town, o,
               PackView[i] (SelfGen[c] town x.Typ x'.Typ, x'.val,
                           unfold (x'.val).vtab.li))
Cast[i, i'] (io, io : ExtView[i]) =
UnpackView[i] (io,
               PackView[i'] (x.Typ, x.val.cobj, x.val.itab.li'))
Prolog[c] (tsubFlds, tsubVtab, (xj : τj)*, •) =
λself : SelfGen[c] OwnFlds[c] tsubFlds tsubVtab. (λxj : τj.)*
let this = PackObj[c] (OwnFlds[c], tsubFlds, tsubVtab, self)
in •
CDecl[c] (classBody : ClassGen[c] OwnFlds[c], •) =
let xc = {town = OwnFlds[c], classBody : ClassGen[c] town}
in •
IDecl[i] (•) =
let xi = {tstamp = AbsMLabels[i],
           Atcobj. λit : ItabGen[i] AbsMLabels[i] tcobj.it
           : ∀tcobj. ItabGen[i] AbsMLabels[i] tcobj
           → ItabGen[i] tstamp tcobj)
in •

```

Figure 11: Term macros used in the translation.

By convention, the type variables used in the translation have the following kinds and intended meaning
(in the context of a template with parameter c or i):

$t_{\text{self}} :: \Omega$	- type of receiver object itself	$t_{\text{twin}} :: \Omega$	- type of an argument/result of receiver's class
$t_{\text{own}} :: \Omega$	- type of private fields of object	$t_{\text{subFds}} :: R^{SCLabels[c]}$	- type of row of field records of a subclass of c
$t_{\text{subVtab}} :: \Omega \rightarrow R^{MLabels[c]}$	- type of row of new methods/interfaces of a subclass of c , parameterized by type of self		
$t_{\text{cobj}} :: \Omega$	- type of object in an interface view	$t_{\text{stamp}} :: R^{MLabels[i]}$	- type of interface's "hidden methods"

Given environments \mathcal{H} and \mathcal{E} ,

$$\begin{aligned} Super[c] &= c' \text{ s.t. } c \prec^C c' \\ SCLabels[c] &= \{l_c' \mid c \leq_{\mathcal{H}}^C c'\} \\ MLabels[c] &= \{l_m \mid \langle m, _ \rangle \in_{\mathcal{H}\mathcal{E}} c\} \cup \{l_i \mid c \ll^C i\} \\ MLabels[i] &= \{l_m \mid \langle m, _ \rangle \in_{\mathcal{H}\mathcal{E}} i\} \cup \{l_{i'} \mid i \leq_{\mathcal{H}}^C i'\} \\ NewPublic[c] &= \lambda t_{\text{twin}}. \lambda t_{\text{subVtab}}. \lambda t_{\text{self}}. ((l_i : ExtItab[i] t_{\text{self}};)^* (l_m : t_{\text{self}} \rightarrow \llbracket T \rrbracket_c t_{\text{twin}};)^* (t_{\text{subVtab}} t_{\text{self}})) \\ &\quad \text{where } i \leftarrow \pi_2(\mathcal{H}(c)), \langle m, T \rangle \leftarrow \pi_2(\mathcal{E}(c)) \end{aligned}$$

$$ObjTmpl[c] = \lambda t_{\text{own}}. \lambda t_{\text{twin}}. \lambda t_{\text{subFds}}. \lambda t_{\text{subVtab}}. \mu t_{\text{self}}. \{ \text{vtab} : VtabGen[c] t_{\text{twin}} t_{\text{subVtab}} t_{\text{self}}; \\ \text{fields: } AllFldsGen[c] t_{\text{own}} t_{\text{subFds}} \}$$

$$ObjGen[c] = \lambda t_{\text{own}}. \mu t_{\text{twin}}. \exists t_{\text{subFds}}. \exists t_{\text{subVtab}}. ObjTmpl[c] t_{\text{own}} t_{\text{twin}} t_{\text{subFds}} t_{\text{subVtab}}$$

$$ExtObj[c] = ObjGen[c] (x_c.\textbf{Typ})$$

$$SelfGen[c] = \lambda t_{\text{own}}. \lambda t_{\text{subFds}}. \lambda t_{\text{subVtab}}. ObjTmpl[c] t_{\text{own}} (ObjGen[c] t_{\text{own}}) t_{\text{subFds}} t_{\text{subVtab}}$$

$$OwnFldsGen[c] = \lambda t_{\text{twin}}. \{(l_f : \text{Ref} (\llbracket T \rrbracket_c t_{\text{twin}}))^*\} \\ \text{where } \langle f, T \rangle \leftarrow \pi_1(\mathcal{E}(c))$$

$$OwnFlds[c] = \mu t_{\text{own}}. OwnFldsGen[c] (ObjGen[c] t_{\text{own}})$$

$$AllFldsGen[c] = \lambda t_{\text{own}}. \lambda t_{\text{subFds}}. AllFlds[Super[c]] (l_c : t_{\text{own}}; t_{\text{subFds}})$$

$$AllFlds[c] = \lambda t_{\text{subFds}}. AllFldsGen[c] (x_c.\textbf{Typ}) t_{\text{subFds}}$$

$$VtabGen[c] = \lambda t_{\text{twin}}. \lambda t_{\text{subVtab}}. \lambda t_{\text{self}}. ExtVtab[Super[c]] (NewPublic[c] t_{\text{twin}} t_{\text{subVtab}}) t_{\text{self}}$$

$$ExtVtab[c] = \lambda t_{\text{subVtab}}. \lambda t_{\text{self}}. VtabGen[c] ExtObj[c] t_{\text{subVtab}} t_{\text{self}} \\ \text{and } ExtVtab[\text{Object}] = \lambda t_{\text{subVtab}}. \lambda t_{\text{self}}. \{t_{\text{subVtab}} t_{\text{self}}\}$$

$$DictGen[c] = \lambda t_{\text{own}}. \lambda t_{\text{subFds}}. \lambda t_{\text{subVtab}}.$$

$$VtabGen[c] (ObjGen[c] t_{\text{own}}) (\lambda t_{\text{self}}. \textbf{Abs}^{MLabels[c]} (SelfGen[c] t_{\text{own}} t_{\text{subFds}} t_{\text{subVtab}}))$$

$$ClassGen[c] = \lambda t_{\text{own}}. \{ \text{dict} : \forall t_{\text{subFds}}. \forall t_{\text{subVtab}}. DictGen[c] t_{\text{own}} t_{\text{subFds}} t_{\text{subVtab}}; \\ \text{init: } \mathbf{1} \rightarrow OwnFldsGen[c] (ObjGen[c] t_{\text{own}}); \\ \text{new: } \mathbf{1} \rightarrow ObjGen[c] t_{\text{own}} \}$$

$$ItabTmpl[i] = \lambda t_{\text{stamp}}. \lambda t_{\text{twin}}. \lambda t_{\text{cobj}}. \{(l_i : ExtItab[i] t_{\text{cobj}};)^* (l_m : t_{\text{cobj}} \rightarrow \llbracket T \rrbracket_i t_{\text{twin}};)^* t_{\text{stamp}}\} \\ \text{where } i' \leftarrow \mathcal{H}(i), \langle m, T \rangle \leftarrow \pi_2(\mathcal{E}(i))$$

$$ViewTmpl[i] = \lambda t_{\text{stamp}}. \lambda t_{\text{twin}}. \lambda t_{\text{cobj}}. \{ \text{cobj} : t_{\text{cobj}}; \text{itab} : ItabTmpl[i] t_{\text{stamp}} t_{\text{twin}} t_{\text{cobj}} \}$$

$$ViewGen[i] = \lambda t_{\text{stamp}}. \mu t_{\text{twin}}. \exists t_{\text{cobj}}. ViewTmpl[i] t_{\text{stamp}} t_{\text{twin}} t_{\text{cobj}}$$

$$ItabGen[i] = \lambda t_{\text{stamp}}. \lambda t_{\text{cobj}}. ItabTmpl[i] t_{\text{stamp}} (ViewGen[i] t_{\text{stamp}}) t_{\text{cobj}}$$

$$ExtView[i] = ViewGen[i] (x_i.\textbf{Typ})$$

$$ExtItab[i] = \lambda t_{\text{cobj}}. ItabGen[i] (x_i.\textbf{Typ}) t_{\text{cobj}}$$

$$\llbracket ty_1 \dots ty_n \rightarrow ty \rrbracket_{ty'} = \lambda t_{\text{twin}} :: \Omega. \llbracket ty_1 \rrbracket_{ty'} t_{\text{twin}} \rightarrow \dots \llbracket ty_n \rrbracket_{ty'} t_{\text{twin}} \rightarrow \llbracket ty \rrbracket_{ty'} t_{\text{twin}}$$

$$\llbracket ty \rrbracket_{ty'} = \lambda t_{\text{twin}} :: \Omega. \begin{cases} t_{\text{twin}}, & \text{if } ty = ty' \\ ExtObj[c] & \text{otherwise, if } ty = c \\ ExtView[i] & \text{otherwise, if } ty = i \end{cases}$$

Figure 10: Type operators used in the translation.

$\frac{(\text{trans-var})}{\begin{array}{c} x \in \text{dom}(\Gamma) \\ ; \vdash \Gamma; \vdash x : \Gamma(x) \rightsquigarrow x \end{array}}$	$\frac{(\text{trans-class})}{\begin{array}{c} \mathcal{H}; \mathcal{E}; \emptyset; c \vdash \exp_j : \text{ty}_j \rightsquigarrow e_j \\ \mathcal{H}; \mathcal{E}; \mathbf{this} : c, \text{ArgTypes}[c](msig_j); c \vdash \exp'_j : \text{RsltType}(msig_j) \rightsquigarrow e'_j \end{array}}$
$\frac{(\text{trans-new})}{\begin{array}{c} \mathcal{H} \vdash_t c' \\ ; \vdash \Gamma; \vdash \mathbf{new} c' : c' \end{array}}$	$\frac{\mathcal{H}; \mathcal{E}; c \vdash \mathbf{new} c' : c'}{\sim \left\{ \begin{array}{l} (x_c \{ \}).\mathbf{new} \{ \}, \text{ if } c' = c \\ x_{c'}.\mathbf{val}.\mathbf{new} \{ \}, \text{ if } c' \neq c \end{array} \right\}}$
$\frac{(\text{trans-get})}{\begin{array}{c} \mathcal{H}; \mathcal{E}; \Gamma; c \vdash \exp : \text{ty}' \rightsquigarrow e \\ \text{ty}' \leq^c_H c \quad \langle f, \text{ty} \rangle \in^c_{\mathcal{E}} c \end{array}}$	$\frac{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash \underline{c \triangleright} \exp.f : \text{ty}}{\sim !((\text{GetFields}[c](e)).l_f)}$
$\frac{(\text{trans-set})}{\begin{array}{c} \mathcal{H}; \mathcal{E}; \Gamma; c \vdash \exp_1 : \text{ty}' \rightsquigarrow e_1 \quad \text{ty}' \leq^c_H c \\ \mathcal{H}; \mathcal{E}; \Gamma; c \vdash \exp_2 : \text{ty} \rightsquigarrow e_2 \quad \langle f, \text{ty} \rangle \in^c_{\mathcal{E}} c \end{array}}$	$\frac{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash \underline{c \triangleright} \exp_1.f = \exp_2 : \text{ty}}{\sim (\text{GetFields}[c](e_1)).l_f := e_2}$
$\frac{(\text{trans-body})}{\begin{array}{c} \forall j \in \{1..n\}. \mathcal{H}; \mathcal{E}; \Gamma; c \vdash \exp_j : \text{ty}_j \rightsquigarrow e_j \\ ; \vdash \Gamma; c \vdash \{ \exp_1; \dots; \exp_n \} : \text{ty}_n \rightsquigarrow (e_1; \dots; e_n) \end{array}}$	$\frac{(\text{trans-class-decl})}{\begin{array}{c} \mathcal{H}'; \mathcal{E}'; \vdash_d \mathbf{class} c \mathbf{ext} c' \mathbf{imp} i^* \{ (ty_j \ f_j = \exp_j;)^* (msig_j \ exp'_j)^* \} \Rightarrow \mathcal{H}; \mathcal{E}; - \\ \mathcal{H}; \mathcal{E}; c; c' \vdash \{ (ty_j \ f_j = \exp_j;)^* (msig_j \ exp'_j)^* \} \rightsquigarrow e \\ \mathcal{H}; \mathcal{E} \vdash p \rightsquigarrow \mathcal{H}''; \mathcal{E}''; e' \end{array}}$
$\frac{(\text{trans-call})}{\begin{array}{c} \mathcal{H}; \mathcal{E}; \Gamma; c \vdash \exp : \text{ty}' \rightsquigarrow e \\ \langle m, (ty_1 \dots ty_n \rightarrow ty) \rangle \in_{\mathcal{H}\mathcal{E}} \text{ty}' \\ \forall j \in \{1..n\}. \mathcal{H}; \mathcal{E}; \Gamma; c \vdash \exp_j : \text{ty}_j \rightsquigarrow e_j \end{array}}$	$\frac{\mathcal{H}'; \mathcal{E}' \vdash \mathbf{class} c \mathbf{ext} c' \mathbf{imp} i^* \{ (ty_j \ f_j = \exp_j;)^* (msig_j \ exp'_j)^* \} p}{\sim \mathcal{H}''; \mathcal{E}''; CDecl[c](Y_v e \{ \}, e')}$
$\frac{(\text{trans-cast})}{\begin{array}{c} \mathcal{H}; \mathcal{E}; \Gamma; c \vdash \exp : \text{ty}' \rightsquigarrow e \quad \text{ty}' \prec_H \text{ty} \\ ; \vdash \Gamma; c \vdash (ty) \exp : \text{ty} \end{array}}$	$\frac{(\text{trans-iface-decl})}{\begin{array}{c} \mathcal{H}'; \mathcal{E}'; \vdash_d \mathbf{interface} i \mathbf{ext} i_1 \dots i_n \{ msig_1; \dots msig_m; \} \Rightarrow \mathcal{H}; \mathcal{E}; - \\ \mathcal{H}; \mathcal{E} \vdash p \rightsquigarrow \mathcal{H}''; \mathcal{E}''; e \end{array}}$
$\frac{\text{where } Own[c](\text{ty}) = \begin{cases} OwnFlds[c], & \text{if } \text{ty} = c \\ x_{\text{ty}}.\mathbf{Typ} & \text{otherwise} \end{cases}}{\sim Cast[\text{ty}', \text{ty}](Own[c](\text{ty}'), Own[c](\text{ty}), e)}$	$\frac{(\text{trans-super})}{\begin{array}{c} c \prec^c_H c' \quad \langle m, (ty_1 \dots ty_n \rightarrow ty) \rangle \in_{\mathcal{H}\mathcal{E}} c' \\ \forall j \in \{1..n\}. \mathcal{H}; \mathcal{E}; \Gamma; c \vdash \exp_j : \text{ty}_j \rightsquigarrow e_j \end{array}}$
	$\frac{\mathcal{H}; \mathcal{E}; \Gamma; c \vdash \underline{c \triangleright} \mathbf{this.super.m}(\exp_1 \dots \exp_n) : \text{ty}}{\sim UnpackObj[c](\mathbf{this}, \mathbf{super}.l_m(x'.\mathbf{val})(e_j)^*)}$

Figure 12: Translation of Javacito to MINIFLINT.

$\lambda_- : \mathbf{1}. \{ \mathbf{x} = \mathbf{ref} 0 \}$. To create the segregated record containing the fields of each superclass, we simply call the `init` functions for each superclass. Suppose the following code is inside the `SPt` package:

```
let init =  $\lambda_- : \mathbf{1}. \{ \mathbf{scale} = \mathbf{ref} 1 \}$ 
  in ... {Pt = Pt.val.init {}}, SPt = init {}} ...
```

Then the type of the record will be `AllFldsGen[SPt]` applied to `OwnFlds[SPt]` and the empty row \mathbf{Abs}^L , where the label set L is $\{Pt, SPt\}$.

5.2 Types used in the translation

5.2.1 Virtual tables

Encoding the type of the vtable of a class is a bit trickier than encoding the fields. We need types for the self argument and for method arguments of the same class. The latter is not simply the same type as external class objects because the private fields need to be accessible.

For now, we will abstract over all of these types, and show how to resolve them in the next section. `VtabGen[c]` produces the type of a vtable for class c . Here is the parameterized type for the vtable of `Pt`:

$$VtabGen[Pt] = \lambda t_{\text{twin}}. \lambda t_{\text{subVtab}}. \lambda t_{\text{self}}. \{ \max : t_{\text{self}} \rightarrow t_{\text{twin}} \rightarrow t_{\text{twin}}; \\ \text{move} : t_{\text{self}} \rightarrow \mathbf{Int} \rightarrow \mathbf{1}; \\ \text{bump} : t_{\text{self}} \rightarrow \mathbf{1}; \\ t_{\text{subVtab}} \ t_{\text{self}} \}$$

where t_{twin} is the type of non-self arguments or results of class `Pt`. t_{subVtab} is the type of a row of new methods in some subclass of `Pt`, parameterized by t_{self} , the type of the receiver object itself.

The tail of the vtable is produced by applying the t_{subVtab} operator to t_{self} . The reason is that new methods added by subclasses will need to use exactly the same self type. To achieve this, we parameterize the row type. The parameterized type of the `SPt` vtable can now be defined using that of `Pt` (we are ignoring the `Zm` interface).

face for now):

$$\begin{aligned} VtabGen[SPt] &= \lambda t_{\text{twin}}. \lambda t_{\text{subVtab}}. \lambda t_{\text{self}}. \\ VtabGen[Pt] \ ExtObj[Pt] &(\lambda t'_{\text{self}}. (\text{zoom} : t'_{\text{self}} \rightarrow \text{Int} \rightarrow \mathbf{1}; \\ &t_{\text{subVtab}} t'_{\text{self}})) t_{\text{self}} \end{aligned}$$

where $\text{ExtObj}[c]$ is the type of an external view of an object of class c , with all private fields hidden. It is instructive to expand the body of this operator, making it independent of the superclass vtable type operator:

$$\begin{aligned} VtabGen[SPt] &= \lambda t_{\text{twin}}. \lambda t_{\text{subVtab}}. \lambda t_{\text{self}}. \{ \\ &\max : t_{\text{self}} \rightarrow \text{ExtObj}[Pt] \rightarrow \text{ExtObj}[Pt]; \\ &\text{move} : t_{\text{self}} \rightarrow \text{Int} \rightarrow \mathbf{1}; \\ &\text{bump} : t_{\text{self}} \rightarrow \mathbf{1}; \\ &\text{zoom} : t_{\text{self}} \rightarrow \text{Int} \rightarrow \mathbf{1}; \\ &t_{\text{subVtab}} t_{\text{self}} \} \end{aligned}$$

In the vtable for class Pt, it is clear that the argument of \max is special because methods in the vtable can access its private fields. In the vtable for class Sp, however, the type of the argument to \max becomes indistinguishable from that of any other Pt object seen from outside, regardless whether we inherit or override \max .

As indicated by the types of functions in the vtable, we are using curried function application to implement methods with arguments. This is only to keep the target calculus as simple as possible. Since we never do partial function application, switching to multi-argument functions would present no problems.

5.2.2 Object types

We combine the field and vtable operators in $\text{ObjTmp}[c]$, a parameterized template for generating class object types (see Figure 10). Note the fixed point for resolving t_{self} ; it ensures that a method selected from the vtab of an object o will accept o as its self argument.

This template may be used to generate several variations of class object types. In order to understand how we resolve the parameters and which variations are needed, we must briefly examine the extensible dictionary for a class:

$$\begin{aligned} \text{dict} &= \Lambda t_{\text{subFlds}}. \Lambda t_{\text{subVtab}}. \\ &\{\max = \lambda \text{self} : (\text{SelfGen}[Pt] \ \text{OwnFlds}[Pt] \ t_{\text{subFlds}} \ t_{\text{subVtab}}), \\ &\quad \lambda \text{other} : (\text{ObjGen}[Pt] \ \text{OwnFlds}[Pt]). \dots, \\ &\quad \dots \} \end{aligned}$$

The difference between $\text{SelfGen}[c]$ and $\text{ObjGen}[c]$ is that the self type should also be extensible; this is why it takes t_{subFlds} and t_{subVtab} as parameters. $\text{ObjGen}[c]$, on the other hand, is an object generated by some unknown class. It might have some additional fields and methods but we cannot know what they are. The $\text{ObjGen}[c]$ operator closes the object type by existentially quantifying the two tail variables and taking a fixed point to resolve t_{twin} (see Figure 10).

The only parameter of $\text{ObjGen}[c]$ is t_{own} , the type of the private fields in c . Inside the class's dictionary, this is instantiated to $\text{OwnFlds}[c]$ so that the private fields are accessible. Outside the class, t_{own} is instantiated instead to $x_c.\text{Typ}$. Thus, the type of a class object viewed from the outside is $\text{ExtObj}[c] = \text{ObjGen}[c] (x_c.\text{Typ})$.

Finally, $\text{SelfGen}[c]$ takes the template and instantiates t_{twin} to $\text{ObjGen}[c] \ t_{\text{own}}$, but leaves the tails as parameters. This ensures that the private fields of self will match those of other arguments of the same class.

5.2.3 Name equivalence

By using abstract types in the fields record of every object type, we encode the class hierarchy and preserve Java's name equivalence. Recall that two abstract types $x.\text{Typ}$ and $x'.\text{Typ}$ are equivalent if and only if x and x' are bound by the same **let**. Thus, structurally equivalent classes will translate to different, incompatible object types.

At first, it might seem that a target-calculus programmer could sabotage name equivalence. Anyone is permitted to use $x_c.\text{val}.\text{init}$ to create records containing values having the abstract types of c and its superclasses. Indeed, a programmer can create a fields record that looks like it belongs to class c , package it with his own devious methods, and have it masquerade as an object of class c . But this is not as subversive as it seems; it is merely subclassing! A Javacito programmer could have achieved the same feat.

5.3 Implementing the dictionary

So far, we have concentrated mostly on typing issues. In this and the next section, we explore the term-level translation of method bodies and expressions. We start with object creation, then move on to method bodies, method invocation, and casts.

5.3.1 Object creation

For the class Pt the extensible dictionary has this shape:

$$\begin{aligned} \text{dict} &= \Lambda t_{\text{subFlds}}. \Lambda t_{\text{subVtab}}. \\ &\{\max = \lambda \text{self}. \dots, \text{move} = \lambda \text{self}. \dots, \text{bump} = \lambda \text{self}. \dots \} \end{aligned}$$

The vtab component of an object is a dictionary *fixed* to the object's type; we obtain it from dict by instantiating to empty rows. To get an object of class Pt we then pair vtab with the fields record (see section 5.1.3), and fold to fix the type of self:

$$\begin{aligned} \text{let } p_0 &= \text{fold } \{ \text{vtab} = \text{dict} [\text{Abs}^{L_c}] [\text{Abs}^{L_m}], \\ &\quad \text{fields} = \{\text{Pt} = \text{init} \{\}\} \} \\ &\text{as } \text{SelfGen}[Pt] \ \text{OwnFlds}[Pt] \ \text{Abs}^{L_c} \ \text{Abs}^{L_m} \end{aligned}$$

where $L_c = \{\text{Pt}\}$ and $L_m = \{\max, \text{move}, \text{bump}\}$.

This object can be passed as the self parameter to methods in the vtab, but its dynamic type is exposed—anyone can see that fields has only one component, Pt. In order to make p_0 's type indistinguishable from the type of objects created by subclasses of Pt, we need to hide the tails:

$$\begin{aligned} \text{let } p_1 &= \text{fold } (t_{\text{subFlds}} = \text{Abs}^{L_c}, \\ &\quad t_{\text{subVtab}} = \text{Abs}^{L_m}, \\ &\quad p_0 : \text{SelfGen}[Pt] \ t_{\text{own}} \ \text{Abs}^{L_c} \ t_{\text{subVtab}}) \\ &\quad : \exists t_{\text{subVtab}}. \text{SelfGen}[Pt] \ t_{\text{own}} \ t_{\text{subFlds}} \ t_{\text{subVtab}} \\ &\text{as } \text{ObjGen}[Pt] \ t_{\text{own}} \end{aligned}$$

where, as usual, t_{own} is $\text{OwnFlds}[Pt]$ if this expression is inside the Pt class package, or $Pt.\text{Typ}$ otherwise. The final **fold** resolves the twin type, as in the non-self argument of \max . This argument has the same private fields accessible as does self, but its additional fields and methods are hidden since it might have a different dynamic type.

This operation, in its general form, is encapsulated as the term macro $\text{PackObj}[c]$ in Figure 11. It is used not only for object creation but for various forms of casting. Its inverse is $\text{UnpackObj}[c]$.

To implement the new function in the class package, we need to wrap a λ around the two **let**-bindings above:

$$\text{let } \text{new} = \lambda _. \mathbf{1}. \text{let } p_0 = \dots \text{ in } \text{let } p_1 = \dots \text{ in } p_1$$

Now we can simply select and apply this function whenever we encounter ‘`new Pt`’, unless it appears in the body of a method in class `Pt` itself. In this case, there is a circularity. The new function refers to `dict`, and a component of `dict` refers to `new`. (Similarly, ‘`new Pt`’ might appear in a field initializer expression, in which case there is a circularity between `init` and `new`.)

We resolve this by passing a *frozen* class record to the call-by-value fixed point combinator Y_v :

```
let  $x_c = \langle t_{\text{own}}::\Omega = \text{OwnFlds}[c],$ 
     $Y_v[\text{ClassGen}[c] \text{ OwnFlds}[c]]$ 
     $(\lambda x_c : \mathbf{1} \rightarrow \text{ClassGen}[c] \text{ OwnFlds}[c].$ 
     $\lambda \_ : \mathbf{1}. \{ \text{dict} = \dots, \text{init} = \dots, \text{new} = \dots \})$ 
   $\{ \}$ 
in ...
```

where `ClassGen[c]` is the type of the class record for class c (see Figure 10).

Using this technique, a method inside `dict` can create a new object of class c with the expression $(x_c \{ \}).\text{new} \{ \}$. Note that the fixed point is *inside* the package, so that the private fields of the new object are exposed. If we applied the fixed point to the whole package, then ‘`(new Pt).x`’ would be illegal.

5.3.2 Method bodies

As demonstrated in Section 5.2.2, methods are translated into curried functions with an explicit `self` parameter. `self` is different from other objects (including twin parameters such as `other`) because it has already been unpacked; it is just a recursive record. To access a field of `other`, one needs to unfold, unpack, unpack again, unfold again, and finally, select the fields record. To access a field of `self`, only the last two steps are necessary.

For convenience, we translate object operations (e.g., method invocation, casts, field selection) uniformly, regardless whether the object is `self`. To support this, we begin each method by packing `self` to make it the same shape as any other object:

```
let this = PackObj[c] (OwnFlds[c],  $t_{\text{subFlds}}$ ,  $t_{\text{subVtab}}$ , self) in ...
```

Now selecting a field from `this` is the same as selecting from `other`. An optimization phase can easily remove any consecutive pack/unpack or fold/unfold expressions.

5.3.3 Method invocation (`invokevirtual`)

To invoke method m on an object o , we just unfold, unpack, unpack again, unfold again (all type manipulations), then select and apply:

```
let  $x = \text{unfold } o$  in let  $x' = x.\text{val}$ 
  in  $\text{unfold } (x'.\text{val}).\text{vtab}.l_m (x'.\text{val})$ 
```

The result can then be applied to any additional arguments.

Since method invocation is not atomic, we must prevent users from selecting a method from one object and applying it to some other object. This error is prevented because the tails of the fields and `vtab` records of o are abstract types. In the code above, $x'.\text{val}.\text{vtab}.l_m$ will have a type which includes the abstract types $x.\text{Typ}$ for the additional fields in o and $x'.\text{Typ}$ for the additional methods in o . Since the only object which can have these types is $x'.\text{val}$ itself, l_m cannot be applied to any other term.

5.3.4 Upward cast

The purpose of an upward cast is to allow an object of class c to masquerade as an object of its superclass, c' . Since we already

keep objects inside nested packages to hide the tails of the fields and `vtab` records, we can simply repackage an object, hiding the components of fields and `vtab` that are part of c but not c' .

In the following example, let o have type `ExtObj[SPt]`:

```
let  $x = \text{unfold } o$ 
in let  $x' = x.\text{val}$ 
  in PackObj[Pt] (Pt.Typ,
    (SPt : SPt.Typ; x.Typ),
    ( $\lambda t_{\text{self}}. (\text{zoom} : t_{\text{self}} \rightarrow \mathbf{Int} \rightarrow \mathbf{1};$ 
       $x'.\text{Typ } t_{\text{self}})),$ 
    x'.val)
```

The result is a value with type `ExtObj[Pt]`. We achieved this using type manipulations only; the underlying object record was not touched. Dynamically, the entire process is a no-op. See the `Cast[c, c']` macro in Figure 11 for the generic form.

5.4 Inheritance and super

Since each class package exports an extensible dictionary, implementing inheritance is straightforward. A class c can select the extensible dictionary of its superclass, and instantiate the tails to its own new fields and methods. For example, inside the `SPt` dict, we bind `super` as follows:

```
dict =  $\Lambda t_{\text{subFlds}}. \Lambda t_{\text{subVtab}}.$ 
let super = Pt.val.dict [SPt : OwnFlds[SPt];  $t_{\text{subFlds}}$ ]
  [ $\lambda t_{\text{self}}. (\text{zoom} : t_{\text{self}} \rightarrow \mathbf{Int} \rightarrow \mathbf{1};$ 
     $t_{\text{subVtab }} t_{\text{self}})]$ 
in {max = super.max,
  move =  $\lambda \text{self} \dots \text{super.move self} \dots,$ 
  bump = super.bump,
  zoom =  $\lambda \text{self} \dots \text{j}$ }
```

Inheritance is implemented simply by selecting methods out of the super dictionary (as in `max` and `bump`, for example) and placing them in the dict of the new subclass. Source language invocations through `super` are implemented by selecting a method from the super dictionary (as in `move`). After the one-time type application to produce `super`, no coercions are necessary.

5.5 Interfaces

Recall our approach to implementing interfaces (see Section 4). An `itable` is an arrangement of some of the methods in a `vtbl`, made to correspond to an interface which the class implements. Given an object of interface type, we know nothing about the shape of its `vtbl`. Thus, we use the `itable` to give a consistent view of the methods in the interface, independent of the underlying `vtbl`.

5.5.1 Interface types

The `itable` is a record of functions, just like the `vtbl`. For the `Zm` interface in our example,

```
 $ItabGen_s[Zm] = \lambda t_{\text{twin}}. \lambda t_{\text{cobj}}. \{ \text{zoom} : t_{\text{cobj}} \rightarrow \mathbf{Int} \rightarrow \mathbf{1} \}$ 
```

where t_{cobj} is the type of the underlying class object (to simplify the presentation, we have omitted the t_{stamp} parameter from `ItabGen`; see section 5.5.4 and figure 10). As before, t_{twin} is the type for non-self arguments of interface `Zm` (there are no such arguments in this particular example).

An interface object (or *view*) is a pair containing some class object (`cobj`) and the appropriate `itable` (`itab`):

```
 $ViewTmpls[Zm] = \lambda t_{\text{twin}}. \lambda t_{\text{cobj}}.$ 
  {cobj : t_{\text{cobj}}; itab : ItabGen_s[Zm] t_{\text{twin }} t_{\text{cobj}}}
```

We resolve the t_{cobj} parameter using an existential type. We should be able to select a method from the itable and pass cobj as its self parameter. We do not know (or it should not matter) what t_{cobj} actually is:

$$\text{View}_s[\text{Zm}] = \mu t_{\text{twin}}. \exists t_{\text{cobj}}. \text{ViewTmpl}_s[\text{Zm}] t_{\text{twin}} t_{\text{cobj}}$$

Finally, we resolve t_{twin} using a recursive type. If the interface Zm had any methods with arguments of type Zm , they would need to look the same, though they might have a different hidden t_{cobj} type.

5.5.2 Invokeinterface

To invoke method m on an interface object io , we bind x to **unfold** io , apply $x.\text{val}.\text{itab}.l_m$ to $x.\text{val}.\text{cobj}$, and apply the result to any additional arguments. Note that this is essentially the same select-and-apply operation that is used for ordinary method invocation (section 5.3.3).

Here again, the abstract type prevents users from doing anything devious with either itab or cobj . The function $x.\text{val}.\text{itab}.l_m$ wants an argument of type $x.\text{Typ}$. The only possible expression of that type is $x.\text{val}.\text{cobj}$.

5.5.3 Interface casts

We have not yet mentioned how to create an interface object from a class object. As discussed in section 4 (and illustrated in Figure 9), the vtable can actually *contain* the itable for all the interfaces a class implements. Then, casting to an interface type is just a matter of selecting the appropriate itable and pairing it with the object itself. For example, the SPt dict would contain the itable for Zm :

```
dict =  $\Lambda t_{\text{subFlds}}. \Lambda t_{\text{subVtab}}.$ 
let super =  $\text{Pt}.\text{val}.\text{dict}[\dots][\dots]$ 
in let newdict =
  {move =  $\lambda \text{self}, \dots \text{super}.\text{move self}, \dots$ ,
   zoom =  $\lambda \text{self}, \dots$ }
  in {max = super.max,
      move = newdict.move,
      bump = super.bump,
      ⇒ Zm = {zoom = newdict.zoom},
      zoom = newdict.zoom}
```

Once this itable is present, the following code can be used to coerce an SPt object o to a Zm interface object:

```
let x = unfold o
in let x' = x.val
  in fold { $t_{\text{cobj}}:\Omega = \text{SelfGen}[\text{SPt}] \text{ SPt}.\text{Typ}$ }  $x.\text{Typ} x'.\text{Typ}$ ,
    { $\text{cobj} = x'.\text{val}$ ,
      $\text{itab} = \text{unfold}(x'.\text{val}).\text{vtab}.\text{Zm}$ }
    :  $\text{ViewTmpl}_s[\text{Zm}] \text{ View}_s[\text{Zm}] t_{\text{cobj}}$ 
  as  $\text{View}_s[\text{Zm}]$ 
```

For this example, we have assumed that this code is contained in some class *other* than SPt . If we were inside SPt , then $\text{SPt}.\text{Typ}$ above would be replaced with $\text{OwnFlds}[\text{SPt}]$. See the $\text{Cast}[c, i]$ macro in Figure 11 for the generic form.

In section 5.3.4, we noted that upward casts consisted of type manipulations only. In contrast, casts to and between interfaces require a simple coercion at runtime. This is not surprising; types can have multiple superinterfaces. A single fixed-offset vtable record cannot, in general, meet the needs of all possible superinterfaces. Either simple coercions or some form of dynamic name lookup are required.

5.5.4 Name equivalence for interfaces

Like class types, interface types in Java use name equivalence. In the formal treatment, we translate interface declarations into a **let**-bound package for the purpose of distinguishing different structurally equivalent interfaces. Each interface package has a hidden *stamp* type ($x_i.\text{Typ}$); the type of an itable for that interface will contain $x_i.\text{Typ}$ as a component. The body of the package ($x_i.\text{val}$) is the identity function, but the argument type accepts structurally equivalent dictionaries and the return type contains the stamp type.

This technique serves not only to distinguish between different, structurally equivalent interface types, but it also, in a sense, prevents forgery. We cannot construct an object that will satisfy interface i unless the interface package x_i is itself accessible.

5.6 Formalization

The formal algorithm translating a (well-typed) *Javacito* program p to a MINIFLINT term is presented in Figure 12, using meta-language macros defined in Figure 11. Some of the macros are parameterized contexts; free variables of a term placed in the hole may be captured in the context. The sequents representing the translation are based on those defining the static semantics of *Javacito* (Figure 5). In the translation of a class declaration we have omitted the straightforward but tedious definition of the macro $\text{MkDict}[c]$, constructing a class dictionary out of a superclass dictionary super , specialized to its subclass c , and a record newdict of the newly defined methods of c . The function of $\text{MkDict}[c]$ is to collect selected components of super (inherited) or newdict (overridden or new) into interface dictionaries and a class dictionary.

The translation is total on type-correct *Javacito* programs, and it maps them to type-correct MINIFLINT terms.

Proposition 3 (Type preservation) *If $\emptyset; \emptyset; \emptyset \vdash_p p : ty \Rightarrow \mathcal{H}; \mathcal{R}$ then $\emptyset; \emptyset \vdash p \rightsquigarrow \mathcal{H}; \mathcal{R}; e$ and $\emptyset \vdash e : [ty]_c \{\}$, where $c \notin \text{dom}(\mathcal{H})$.*

The semantic correctness of the translation is shown by establishing a mapping from *Javacito* runtime configurations to MINIFLINT runtime configurations which maps value-configurations to value-configurations, and proving that a *Javacito* reduction step corresponds with respect to this mapping to finitely many MINIFLINT reduction steps [23]; we omit the proofs due to space constraints.

6 Extensions

Many features of Java can be accommodated directly in our framework, but were left out in order to simplify the formal presentation.

6.1 Unordered records

We have used ordered records in our translation of *Javacito*, but we can add unordered records [38] into the target language as well:

$$\begin{array}{ll} \text{types} & \tau ::= \dots | \{\!\!\{\tau\}\!\!\} \\ \text{terms} & e ::= \dots | \{\!\!\{(l = e)^*\}\!\!\} \end{array}$$

Here, $\{\!\!\{\tau\}\!\!\}$ denotes the unordered record type where τ must be of kind R^\emptyset , and $\{\!\!\{(l = e)^*\}\!\!\}$ is the unordered record term which (unlike ordered records) requires runtime dictionary construction. Selecting from an unordered record uses the same syntax as selecting from an ordered record but it requires runtime dictionary lookup.

Some Java compilers use the same representation for objects of interface types as for those of class types, so casting to interface

type requires no runtime operations. Unordered records can be used to support this representation by collecting all the itable entries of a vtable into a separate unordered record, itself an element of the still ordered vtable. Casting an object into interface type only requires repackaging it (a runtime no-op) to hide those entries not exported by the current interface.

6.2 Access control

Javacito only allows private fields and public methods, but other access scoping schemes can also be supported. Private methods can simply be `let`-bound within the class dictionary since they can neither be called from subclasses nor overridden. Public fields could be placed directly in an object’s fields record, without needing to be segregated.

Protected and package scopes require adding a notion of Java-style *package* into MINIFLINT. Otherwise, they can be supported using the technique proposed by Moby [16]. The `dict` and `new` fields in our class encoding roughly correspond to the class view and the object view in Moby. If we export a class outside its definitional package, all protected methods and fields should be hidden from the object view but not the class view while those of package scope should be hidden from both.

6.3 Dynamic casts and exception

Dynamic casts and runtime type identification can be addressed using extensible variant types similar to the `exn` data type in Standard ML. We add a new extensible variant declaration inside each class c or interface i (for simplicity we borrow the syntax of exception declaration in ML):

```
exception Tagc of ObjGen[c] town;
or
exception Tagi of ViewGen[i] tstamp;
```

Suppose our extensible variant type is named `Tagged`; then, inside each class declaration of c , we add a method named `allviews` that takes the self object and returns a list of `Tagged` values. More specifically, the `allviews` method takes the self object s , and for each ancestor class c' (and interface i') it upward-casts s into an object of c' (or i'), packages it with the $Tag_{c'}$ (or $Tag_{i'}$) flag, and then returns the entire list of resulting `Tagged` values. We also add the specification for the `allviews` method into each interface so that it can be invoked on objects of interface type as well.

To test if an object o is an `instanceof` class c (or interface i), we invoke the `allviews` method on o and then check through the result `Tagged` list to see if any of them is tagged with Tag_c (or Tag_i). The `checkcast` operation can be implemented in the same way since the object associated with Tag_c (or Tag_i) is already of type c (or i). Notice unlike Glew’s recent work [19], our technique does not use hierarchical extensible sums or variance-based subtyping so it requires a much simpler target language. With careful coding, the `allviews` method can be implemented as efficiently as the untyped code used in typical Java compilers.

6.4 Miscellaneous

Null references are encoded by lifting all external object types to variant types with a `null` alternative, similar to the `option` data type in Standard ML. Then, all object operations must first verify that the object pointer is not `null`. Since this could fail, we also need support for unchecked exceptions, which can be achieved by adding an exception mechanism in the style of Standard ML.

Mutually recursive class and interface declarations need a fixed-point construction over the corresponding components in

MINIFLINT; the fixed point already used in the definition of a single class cannot be generalized for this purpose since this will result in “friend” status of the other classes.

Other aspects of Java such as concurrency, `final`, reflection, class loaders, and dynamic loading are more challenging; we have preliminary ideas about implementing them, but developing the details remains as future work. Of course our F^ω -based framework is particularly well-suited for implementing Java extensions based on higher-order functions and types such as those in Pizza [27].

7 Related Work

Object and class encodings have been extensively studied. The use of row polymorphism positions our scheme most closely to that of Rémy and Vouillon [32], however the special object types of Objective ML reduce the use of row polymorphism to only the cases of binary methods [4], while the self-application semantics of our scheme uses rows to represent the open type of self even though Java lacks binary methods. The unordered records of Objective ML are geared towards support for multiple inheritance; in contrast our scheme makes single inheritance efficient, while an extension with multiple inheritance is still possible (but less efficient) using interface views as the basic representation.

In the context of object encodings in F^ω -based languages [5, 29], the genealogy of our encoding can be traced back to encodings based on F-bounded polymorphism [7, 13], from which the F-bounds have been eliminated using intersection types, which have further been replaced by row polymorphism (in the case of objects) or realized as tuples (in the case of interface views). At the most basic level, method invocation uses self-application (the whole object is a parameter of each method), showing similarity with the encoding due to Abadi, Cardelli, and Viswanathan [1]; however, hiding the actual class of the receiver is achieved using existential quantification over row variables instead of splitting the object into a known interface and a hidden implementation. This allows reuse of methods in subclasses without any overhead (modulo type applications, which are no-ops in the intended type-erasure semantics). Further we use an analog of the recursive-existential encoding due to Bruce [6] to give types to other arguments or results belonging to the same class or a subclass, as needed in Java, without over-restricting the type to be the same as the receiver’s. The private instance variables of these objects are thus accessible by methods of the class; they are protected by the final level of encapsulation.

Fisher and Mitchell [14, 15] show how to use extensible objects to model Java-like class constructs. Our encoding does not rely on extensible objects as primitives, but it may be viewed as an implementation of some of their properties in terms of simpler constructs. In particular, extensibility of objects is only used when they have the role of prototypes, and the ability to extend an object in general takes us further from the intended Java semantics. In our translation these roles are clearly separated. Our encoding of classes guarantees restricted visibility and (as a consequence) correct initialization of private instance variables, and they provide directly reusable collections of methods as well as means to create new objects. One criterion our encoding fails is to automatically propagate changes in a base class to its descendants, but it is unclear if this criterion can coexist with Java’s binary compatibility.

8 Conclusions

We have presented a formal translation of Java classes, interfaces, and privacy into a call-by-value variant of F^ω using simple and well-known extensions. Even though the resulting code contains full type information, the runtime object layout corresponds to what

one might expect in an untyped implementation. The operations of object creation, method invocation, and field selection are implemented efficiently in terms of primitive F^ω constructs. Thus, they are candidates for standard optimizations and we can reason about their interaction with foreign code.

An implementation of this encoding is in progress. As of April 1999, we finished a prototype implementation which can translate all the features of *Javacito* into a simple, interpreted target calculus. We are working on connecting this simple target calculus implementation to the full-fledged FLINT compiler [34, 35], in order to leverage its type-directed optimizations, compiler back ends, and runtime support. The actual FLINT intermediate language is surprisingly close to the target calculus presented here (MINIFLINT); all we need is just to add the row polymorphism.

Acknowledgments

We wish to thank the anonymous referees for their suggestions.

References

- [1] M. Abadi, L. Cardelli, and R. Viswanathan. An interpretation of objects and object types. In *Proc. 23rd ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages (POPL'96)*, pages 396–409, St. Petersburg, January 1996.
- [2] A. W. Appel and D. B. MacQueen. Standard ML of New Jersey. In M. Wirsing, editor, *Third Int'l Symp. on Prog. Lang. Implementation and Logic Programming*, pages 1–13, New York, August 1991. Springer-Verlag.
- [3] N. Benton, A. Kennedy, and G. Russell. Compiling Standard ML to Java bytecodes. In *Proc. 1998 Int'l Conf. on Functional Programming (ICFP'98)*, pages 129–140, September 1998.
- [4] K. Bruce, L. Cardelli, G. Castagna, The Hopkins Objects Group, G. T. Leavens, and B. Pierce. On binary methods. *Theory and Practice of Object Systems*, 1(3):221–242, December 1995.
- [5] K. Bruce, L. Cardelli, and B. Pierce. Comparing object encodings. *Information and Computation (to appear)*, 1998.
- [6] K. B. Bruce. A paradigmatic object-oriented programming language: Design, static typing and semantics. *Journal of Functional Programming*, 4(2), 1994.
- [7] P. Canning, W. Cook, W. Hill, J. Mitchell, and W. Olthoff. F-bounded polymorphism for object-oriented programming. In *Proceedings of the Conference on Functional Programming Languages and Computer Architecture*, pages 273–280, 1989.
- [8] L. Cardelli and X. Leroy. Abstract types and the dot notation. In *Proc. IFIP Working Conf. on Programming Concepts and Methods*, pages 466–491, Israel, April 1990.
- [9] L. Cardelli and G. Longo. A semantic basis for Quest. *Journal of Functional Programming*, 1(4), 1991.
- [10] J. Dean, G. DeFouw, D. Grove, V. Litvinov, and C. Chambers. VorTEX: An optimizing compiler for object-oriented languages. In *Proc. ACM SIGPLAN '96 Conf. on Object-Oriented Programming Systems, Languages, and applications*, pages 83–100, October 1996.
- [11] S. Drossopoulou and S. Eisenbach. Java is type safe—probably. In *Proc. 11th European Conf. on Object-Oriented Prog.*, February 1997.
- [12] S. Drossopoulou and S. Eisenbach. Towards an operational semantics and proof of type soundness for Java, April 1998.
- [13] J. Eifrig, S. Smith, V. Trifonov, and A. Zwarico. An interpretation of typed OOP in a language with state. *Lisp and Symbolic Computation*, 8(4):357–397, 1995.
- [14] K. Fisher. *Type Systems for Object-Oriented Programming Languages*. PhD thesis, Stanford University, August 1996.
- [15] K. Fisher and J. Mitchell. On the relationship between classes, objects, and data abstraction. *Theory and Practice of Object Systems*, 4(1):3–25, January 1998.
- [16] K. Fisher and J. Reppy. The design of a class mechanism for Moby. In *Proc. ACM SIGPLAN '99 Conf. on Prog. Lang. Design and Implementation*, pages 37–49, May 1999.
- [17] M. Flatt, S. Krishnamurthi, and M. Felleisen. A programmer's reduction semantics for classes and mixins. Technical Report 97-293, Rice University, 1997.
- [18] J. Y. Girard. *Interprétation Fonctionnelle et Elimination des Coupures dans l'Arithmétique d'Ordre Supérieur*. PhD thesis, University of Paris VII, 1972.
- [19] N. Glew. Type dispatch for named hierarchical types. In *Proc. 1999 ACM SIGPLAN International Conference on Functional Programming (ICFP'99)*, September 1999.
- [20] N. Heintze and J. G. Riecke. The slam calculus: Programming with secrecy and integrity. In *Twenty-fifth Annual ACM Symp. on Principles of Prog. Languages*, pages 365–377, Jan 1998.
- [21] M. Hofmann and B. Pierce. A unifying type-theoretic framework for objects. *Journal of Functional Programming*, January 1994.
- [22] A. Krall and R. Grafl. CACAO—a 64-bit Java VM Just-In-Time compiler. In *Proc. ACM PPoPP'97 Workshop on Java for Science and Engineering Computation.*, 1997.
- [23] C. League, Z. Shao, and V. Trifonov. Representing Java classes in a typed intermediate language (extended version). Technical Report YALEU/DCS/RR-1180, Department of Computer Science, Yale University, New Haven, CT, June 1999.
- [24] X. Leroy and F. Rouaix. Security properties of typed applets. In *Twenty-fifth Annual ACM Symp. on Principles of Prog. Languages*, pages 391–403, Jan 1998.
- [25] T. Lindholm and F. Yellin. *The Java Virtual Machine Specification (Second Edition)*. Addison-Wesley, 1999.
- [26] Microsoft Win32 VM for Java object model, 1998.
- [27] M. Odersky and P. Wadler. Pizza into Java: Translating theory into practice. In *Twenty-fourth Annual ACM Symp. on Principles of Prog. Languages*, pages 146–159, Jan 1997.
- [28] J. Palsberg and P. Ørbæk. Trust in the λ -calculus. *Journal of Functional Programming*, 7(6):557–591, November 1997.
- [29] B. C. Pierce and D. N. Turner. Simple type-theoretic foundations for object-oriented prog. *Journal of Functional Programming*, 4(2):207–247, April 1994.
- [30] T. A. Proebsting, G. Townsend, P. Bridges, J. H. Hartman, T. Newsham, and S. A. Watterson. Toba: Java for applications: A way ahead of time (WAT) compiler. In *Proc. Third Conf. on Object-Oriented Technologies and Systems (COOTS'97)*, 1997.
- [31] D. Rémy. Syntactic theories and the algebra of record terms. Technical Report 1869, INRIA, 1993.
- [32] D. Rémy and J. Vouillon. Objective ML: A simple object-oriented extension of ML. In *Twenty-fourth Annual ACM Symp. on Principles of Prog. Languages*, pages 40–53, Jan 1997.
- [33] J. C. Reynolds. Towards a theory of type structure. In *Proc., Colloque sur la Programmation, Lecture Notes in Computer Science, volume 19*, pages 408–425. Springer-Verlag, Berlin, 1974.
- [34] Z. Shao. An overview of the FLINT/ML compiler. In *Proc. 1997 ACM SIGPLAN Workshop on Types in Compilation*, June 1997.
- [35] Z. Shao. Typed common intermediate format. In *Proc. USENIX Conf. on Domain-Specific Languages*, pages 89–101, Santa Barbara, CA, October 1997.
- [36] Z. Shao. Typed cross-module compilation. In *Proc. 1998 ACM SIGPLAN International Conference on Functional Programming (ICFP'98)*, pages 141–152, September 1998.
- [37] Z. Shao, C. League, and S. Monnier. Implementing typed intermediate languages. In *Proc. 1998 ACM SIGPLAN International Conference on Functional Programming (ICFP'98)*, pages 313–323, September 1998.
- [38] A. Wright, S. Jaganathan, C. Ungureanu, and A. Hertzmann. Compiling Java to a typed lambda-calculus: A preliminary report. In X. Leroy and A. Ohori, editors, *Proc. 1998 International Workshop on Types in Compilation: LNCS Vol 1473*, pages 9–27, Kyoto, Japan, March 1998. Springer-Verlag.